

EXHIBIT C

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

NVIDIA CORPORATION,
Petitioners,

v.

ADVANCED CLUSTER SYSTEMS, INC.
Patent Owner

Case No. IPR2021-00019
U.S. Patent 10,333,768

DECLARATION OF DEAN DAUGER, Ph.D.

IPR2021-00019

NVIDIA v. Advanced Cluster Systems

provide customers with Internet-based access to instances of gridMathematica running on ACS's cluster. Second, ACS was investigating whether it could make its cluster available to its customers to run their own parallel applications on the cluster, while the cluster was not otherwise busy. ACS purchased a license to the then-current version of Pooch, Pooch Pro, from Dauger Research to allow ACS's customers to access the cluster over the Internet, initialize the cluster with their own parallel applications, and then run their applications over the cluster. ACS retained Dauger Research to help integrate Pooch Pro into the ACS cluster and to help with the MGT product.

22. Through the course of my work for ACS, Mr. Tannenbaum and I realized certain limitations of the MGT product due to the underlying gridMathematica program. GridMathematica was a mathematics program developed by Wolfram Research. GridMathematica—as its name would suggest—was designed on the grid computing model. Within the “grid” there was one master gridMathematica node and several slave gridMathematica nodes. The slave gridMathematica nodes could not communicate with each other in the peer-to-peer and collective fashion, used in cluster computing. Each slave node could only communicate with the master gridMathematica node, which would then forward the instruction or data to the appropriate gridMathematica slave node. Because of this grid computing architecture, gridMathematica was ill-suited for problems with a

IPR2021-00019

NVIDIA v. Advanced Cluster Systems

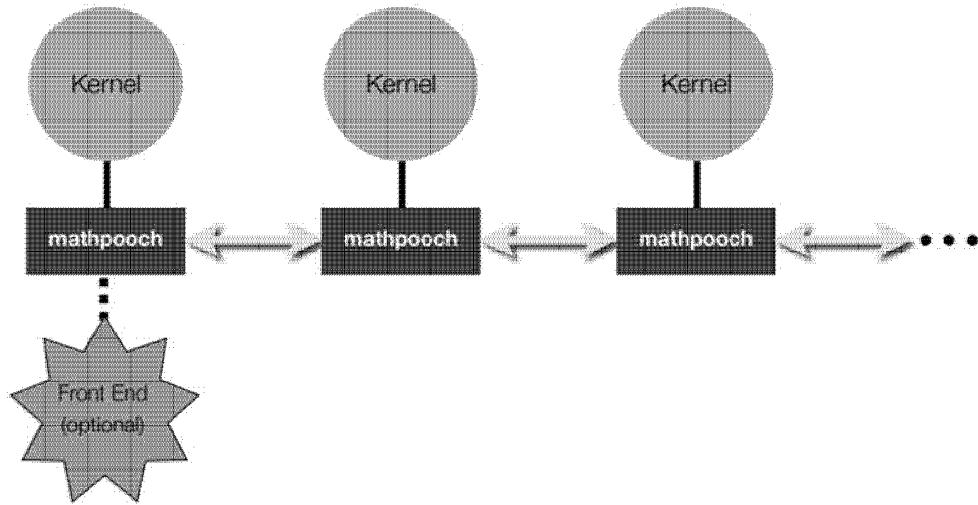
high degree of dependency. To get the highest possible parallel computing performance required cluster computing.

23. In view of these shortcomings, Mr. Tannenbaum and I began to explore ways to improve MGT by adding the peer-to-peer and collective internode communication needed to optimize parallel computation of Mathematica on a cluster. Over time our collaboration determined that we could do so using a new cluster-computing architecture. The new architecture interposed a communication layer between the Mathematica front end and the Mathematica kernels running on each node of the cluster, or back end. This communication layer enabled asynchronous peer-to-peer communication between each node of the cluster, creating a cluster implementation of Mathematica. ACS leveraged the model-view-controller architecture in modern applications like Mathematica. ACS later named the ultimate commercial product SEM.¹ Below is a diagram from one of the first SEM manuals, which illustrates this architecture:

¹ Mathpooch, PoochMPI, semath, and SEM are different names for the same parallelization framework. Similarly, Mathpooch module, PoochMPI module, semath module, and SEM module are different names for particular modules of Mathpooch, PoochMPI, semath, and SEM on a given node. For clarity this declaration will refer to them as SEM/SEM modules, although references to Mathpooch, PoochMPI, and semath also refer to SEM.

IPR2021-00019

NVIDIA v. Advanced Cluster Systems



Ex. 2011 (SEM Manual) at 4. As the SEM manual noted, the network topology between the SEM module on each node could be all-to-all, not just nearest neighbor.

24. SEM's architecture allowed it to break the performance/programming barrier. SEM interposes a communication layer between the front end user interface (view) and the back end Mathematica kernels (model) on each node, and that communication layer manages the peer-to-peer behavior of the nodes. As a result of this architecture, the user is minimally exposed to the need for low-level MPI calls. By reducing the user's exposure to MPI programming, SEM's architecture significantly reduces the difficulty of parallel programming.

25. After testing SEM and determining that the new architecture worked surprisingly well, Mr. Tannenbaum asked me to proceed to generalize the concept for use with other applications beyond Mathematica. We developed SET, which

IPR2021-00019

NVIDIA v. Advanced Cluster Systems

76. Attached to this Patent Owner's Preliminary Response as Exhibit 2025 is a true and accurate copy of an Excerpt from Zvi Tannenbaum, *SETTM Windows Project Narrative: Topic 2b*.

77. Attached to this Patent Owner's Preliminary Response as Exhibit 2026 is a true and accurate copy of an Excerpt from U.S. Department of Energy, Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) Program, *Topics: FY 2015 Phase I Release 1* (Ver. 1 July 14, 2014).

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Date: 2/9/2021



Dean Dauger, Ph.D.

34386903